

A LOW-RANK MULTIGRID METHOD FOR THE STOCHASTIC STEADY-STATE DIFFUSION PROBLEM*

HOWARD C. ELMAN[†] AND TENGFEI SU[‡]

Abstract. We study a multigrid method for solving large linear systems of equations with tensor product structure. Such systems are obtained from stochastic finite element discretization of stochastic partial differential equations such as the steady-state diffusion problem with random coefficients. When the variance in the problem is not too large, the solution can be well approximated by a low-rank object. In the proposed multigrid algorithm, the matrix iterates are truncated to low rank to reduce memory requirements and computational effort. The method is proved convergent with an analytic error bound. Numerical experiments show its effectiveness in solving the Galerkin systems compared to the original multigrid solver, especially when the number of degrees of freedom associated with the spatial discretization is large.

Key words. stochastic finite element method, multigrid, low-rank approximation

AMS subject classifications. 35R60, 60H15, 60H35, 65F10, 65N30, 65N55

1. Introduction. Stochastic partial differential equations (SPDEs) arise from physical applications where the parameters of the problem are subject to uncertainty. Discretization of SPDEs gives rise to large linear systems of equations which are computationally expensive to solve. These systems are in general sparse and structured. In particular, the coefficient matrix can often be expressed as a sum of tensor products of smaller matrices [6, 13, 14]. For such systems it is natural to use an iterative solver where the coefficient matrix is never explicitly formed and matrix-vector products are computed efficiently. One way to further reduce costs is to construct low-rank approximations to the desired solution. The iterates are truncated so that the solution method handles only low-rank objects in each iteration. This idea has been used to reduce the costs of iterative solution algorithms based on Krylov subspaces. For example, a low-rank conjugate gradient method was given in [9], and low-rank generalized minimal residual methods have been studied in [2, 10].

In this study, we propose a low-rank multigrid method for solving the Galerkin systems. We consider a steady-state diffusion equation with random diffusion coefficient as model problem, and we use the stochastic finite element method (SFEM, see [1, 7]) for the discretization of the problem. The resulting Galerkin system has tensor product structure and moreover, quantities used in the computation, such as the solution sought, can be expressed in matrix format. It has been shown that such systems admit low-rank approximate solutions [3, 9]. In our proposed multigrid solver, the matrix iterates are truncated to have low rank in each iteration. We derive an analytic bound for the error of the solution and show the convergence of the algorithm. We demonstrate using benchmark problems that the low-rank multigrid solver is often more efficient than a solver that does not use truncation, and that it is especially advantageous in reducing computing time for large-scale problems.

An outline of the paper is as follows. In section 2 we state the problem and

*This work was supported by the U.S. Department of Energy Office of Advanced Scientific Computing Research, Applied Mathematics program under award DE-SC0009301 and by the U.S. National Science Foundation under grant DMS1418754.

[†]Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 (elman@cs.umd.edu).

[‡]Applied Mathematics & Statistics, and Scientific Computation Program, University of Maryland, College Park, MD 20742 (tengfesu@math.umd.edu).

briefly review the stochastic finite element method and the multigrid solver for the stochastic Galerkin system from which the new technique is derived. In section 3 we discuss the idea of low-rank approximation and introduce the multigrid solver with low-rank truncation. A convergence analysis of the low-rank multigrid solver is also given in this section. The results of numerical experiments are shown in section 4 to test the performance of the algorithm, and some conclusions are drawn in the last section.

2. Model problem. Consider the stochastic steady-state diffusion equation with homogeneous Dirichlet boundary conditions

$$(2.1) \quad \begin{cases} -\nabla \cdot (c(x, \omega) \nabla u(x, \omega)) = f(x) & \text{in } D \times \Omega, \\ u(x, \omega) = 0 & \text{on } \partial D \times \Omega. \end{cases}$$

Here D is a spatial domain and Ω is a sample space with σ -algebra \mathcal{F} and probability measure P . The diffusion coefficient $c(x, \omega) : D \times \Omega \rightarrow \mathbb{R}$ is a random field. We consider the case where the source term f is deterministic. The stochastic Galerkin formulation of (2.1) uses a weak formulation: find $u(x, \omega) \in \mathbb{V} = H_0^1(D) \otimes L^2(\Omega)$ satisfying

$$(2.2) \quad \int_{\Omega} \int_D c(x, \omega) \nabla u(x, \omega) \cdot \nabla v(x, \omega) dx dP = \int_{\Omega} \int_D f(x) v(x, \omega) dx dP$$

for all $v(x, \omega) \in \mathbb{V}$. The problem is well posed if $c(x, \omega)$ is bounded and strictly positive, i.e.,

$$0 < c_1 \leq c(x, \omega) \leq c_2 < \infty, \text{ a.e. } \forall x \in D,$$

so that the Lax-Milgram lemma establishes existence and uniqueness of the weak solution.

We will assume that the stochastic coefficient $c(x, \omega)$ is represented as a truncated Karhunen-Loève (KL) expansion [11, 12], in terms of a finite collection of uncorrelated random variables $\{\xi_l\}_{l=1}^m$:

$$(2.3) \quad c(x, \omega) \approx c_0(x) + \sum_{l=1}^m \sqrt{\lambda_l} c_l(x) \xi_l(\omega)$$

where $c_0(x)$ is the mean function, $(\lambda_l, c_l(x))$ is the l th eigenpair of the covariance function $r(x, y)$, and the eigenvalues $\{\lambda_l\}$ are assumed to be in non-increasing order. In section 4 we will further assume these random variables are independent and identically distributed. Let $\rho(\xi)$ be the joint density function and Γ be the joint image of $\{\xi_l\}_{l=1}^m$. The weak form of (2.1) is then given as follows: find $u(x, \xi) \in \mathbb{W} = H_0^1(D) \otimes L^2(\Gamma)$ s.t.

$$(2.4) \quad \int_{\Gamma} \rho(\xi) \int_D c(x, \xi) \nabla u(x, \xi) \cdot \nabla v(x, \xi) dx d\xi = \int_{\Gamma} \rho(\xi) \int_D f(x) v(x, \xi) dx d\xi$$

for all $v(x, \xi) \in \mathbb{W}$.

2.1. Stochastic finite element method. We briefly review the stochastic finite element method as described in [1, 7]. This method approximates the weak solution of (2.1) in a finite-dimensional subspace

$$(2.5) \quad \mathbb{W}^{hp} = S^h \otimes T^p = \text{span}\{\phi(x)\psi(\xi) \mid \phi(x) \in S^h, \psi(\xi) \in T^p\},$$

where S^h and T^p are finite-dimensional subspaces of $H_0^1(D)$ and $L^2(\Gamma)$. We will use quadrilateral elements and piecewise bilinear basis functions $\{\phi(x)\}$ for the discretization of the physical space $H_0^1(D)$, and generalized polynomial chaos [17] for the stochastic basis functions $\{\psi(\xi)\}$. The latter are m -dimensional orthogonal polynomials whose total degree doesn't exceed p . The orthogonality relation means

$$\int_{\Gamma} \psi_r(\xi) \psi_s(\xi) \rho(\xi) d\xi = \delta_{rs} \int_{\Gamma} \psi_r^2(\xi) \rho(\xi) d\xi.$$

For instance, Legendre polynomials are used if the random variables have uniform distribution with zero mean and unit variance. The number of degrees of freedom in T^p is

$$N_{\xi} = \frac{(m+p)!}{m!p!}.$$

Given the subspace, now one can write the SFEM solution as a linear combination of the basis functions,

$$(2.6) \quad u_{hp}(x, \xi) = \sum_{j=1}^{N_x} \sum_{s=1}^{N_{\xi}} u_{js} \phi_j(x) \psi_s(\xi),$$

where N_x is the dimension of the subspace S^h . Substituting (2.3) and (2.6) into (2.4), and taking the test function as any basis function $\phi_i(x) \psi_r(\xi)$ results in the Galerkin system: find $\mathbf{u} \in \mathbb{R}^{N_x N_{\xi}}$, s.t.

$$(2.7) \quad A \mathbf{u} = \mathbf{f}.$$

The coefficient matrix A can be represented in tensor product notation [14],

$$(2.8) \quad A = G_0 \otimes K_0 + \sum_{l=1}^m G_l \otimes K_l,$$

where $\{K_l\}_{l=0}^m$ are the stiffness matrices and $\{G_l\}_{l=0}^m$ correspond to the stochastic part, with entries

$$(2.9) \quad \begin{aligned} G_0(r, s) &= \int_{\Gamma} \psi_r(\xi) \psi_s(\xi) \rho(\xi) d\xi, \quad K_0(i, j) = \int_D c_0(x) \nabla \phi_i(x) \nabla \phi_j(x) dx, \\ G_l(r, s) &= \int_{\Gamma} \xi_l \psi_r(\xi) \psi_s(\xi) \rho(\xi) d\xi, \quad K_l(i, j) = \int_D \sqrt{\lambda_l} c_l(x) \nabla \phi_i(x) \nabla \phi_j(x) dx, \end{aligned}$$

$l = 1, \dots, m$; $r, s = 1, \dots, N_{\xi}$; $i, j = 1, \dots, N_x$. The right-hand side can be written as a tensor product of two vectors:

$$(2.10) \quad \mathbf{f} = g_0 \otimes f_0,$$

where

$$(2.11) \quad \begin{aligned} g_0(r) &= \int_{\Gamma} \psi_r(\xi) \rho(\xi) d\xi, \quad r = 1, \dots, N_{\xi}, \\ f_0(i) &= \int_D f(x) \phi_i(x) dx, \quad i = 1, \dots, N_x. \end{aligned}$$

Note that in the Galerkin system (2.7), the matrix A is symmetric and positive definite. It is also blockwise sparse (see Figure 2.1) due to the orthogonality of $\{\psi_r(\xi)\}$.

The size of the linear system is in general very large ($N_x N_\xi \times N_x N_\xi$). For such a system it is suitable to use an iterative solver. Multigrid methods are among the most effective iterative solvers for the solution of discretized elliptic PDEs, capable of achieving convergence rates that are independent of the mesh size, with computational work growing only linearly with the problem size [8, 15].




FIG. 2.1. Block structure of A . $m = 4, p = 1, 2, 3$ from left to right. Block size is $N_x \times N_x$.

2.2. Multigrid. In this subsection we discuss a geometric multigrid solver proposed in [4] for the solution of the stochastic Galerkin system (2.7). For this method, the mesh size h varies for different grid levels, while the polynomial degree p is held constant, i.e., the fine grid space and coarse grid space are defined as

$$(2.12) \quad \mathbb{W}^{hp} = S^h \otimes T^p, \quad \mathbb{W}^{2h,p} = S^{2h} \otimes T^p,$$

respectively. Then the prolongation and restriction operators are of the form

$$(2.13) \quad \mathcal{P} = I \otimes P, \quad \mathcal{R} = I \otimes P^T,$$

where P is the same prolongation matrix as in the deterministic case. On the coarse grid we only need to construct matrices $\{K_l^{2h}\}_{l=0}^m$, and

$$(2.14) \quad A^{2h} = G_0 \otimes K_0^{2h} + \sum_{l=1}^m G_l \otimes K_l^{2h}.$$

The matrices $\{G_l\}_{l=0}^m$ are the same for all grid levels.

Algorithm 2.1 describes the complete multigrid method. In each iteration, we apply one multigrid cycle (VCYCLE) for the residual equation

$$(2.15) \quad A\mathbf{c}^{(i)} = \mathbf{r}^{(i)} = \mathbf{f} - A\mathbf{u}^{(i)}$$

and update the solution $\mathbf{u}^{(i)}$ and residual $\mathbf{r}^{(i)}$. The VCYCLE function is called recursively. On the coarsest grid level ($h = h_0$) we form matrix A and solve the linear system directly. The system is of order $O(N_\xi)$ since $A \in \mathbb{R}^{N_x N_\xi \times N_x N_\xi}$ where N_x is a very small number on the coarsest grid. The smoothing function (SMOOTH) is based on a matrix splitting $A = Q - Z$ and stationary iteration

$$(2.16) \quad \mathbf{u}_{s+1} = \mathbf{u}_s + Q^{-1}(\mathbf{f} - A\mathbf{u}_s),$$

which we assume is convergent, i.e., the spectral radius $\rho(I - Q^{-1}A) < 1$. The algorithm is run until the specified relative tolerance tol or maximum number of iterations $maxit$ is reached. It is shown in [4] that for $f \in L^2(D)$, the convergence rate of this algorithm is independent of the mesh size h , the number of random variables m , and the polynomial degree p .

Algorithm 2.1: Multigrid for stochastic Galerkin systems

```

1: initialization:  $i = 0$ ,  $\mathbf{r}^{(0)} = \mathbf{f}$ ,  $r_0 = \|\mathbf{f}\|_2$ 
2: while  $r > tol * r_0$  &  $i \leq maxit$  do
3:    $\mathbf{c}^{(i)} = \text{VCYCLE}(A, \mathbf{0}, \mathbf{r}^{(i)})$ 
4:    $\mathbf{u}^{(i+1)} = \mathbf{u}^{(i)} + \mathbf{c}^{(i)}$ 
5:    $\mathbf{r}^{(i+1)} = \mathbf{f} - A\mathbf{u}^{(i+1)}$ 
6:    $r = \|\mathbf{r}^{(i+1)}\|_2$ ,  $i = i + 1$ 
7: end

8: function  $\mathbf{u}^h = \text{VCYCLE}(A^h, \mathbf{u}_0^h, \mathbf{f}^h)$ 
9:   if  $h == h_0$  then
10:     solve  $A^h \mathbf{u}^h = \mathbf{f}^h$  directly
11:   else
12:      $\mathbf{u}^h = \text{SMOOTH}(A^h, \mathbf{u}_0^h, \mathbf{f}^h)$ 
13:      $\mathbf{r}^h = \mathbf{f}^h - A^h \mathbf{u}^h$ 
14:      $\mathbf{r}^{2h} = \mathcal{R} \mathbf{r}^h$ 
15:      $\mathbf{c}^{2h} = \text{VCYCLE}(A^{2h}, \mathbf{0}, \mathbf{r}^{2h})$ 
16:      $\mathbf{u}^h = \mathbf{u}^h + \mathcal{P} \mathbf{c}^{2h}$ 
17:      $\mathbf{u}^h = \text{SMOOTH}(A^h, \mathbf{u}^h, \mathbf{f}^h)$ 
18:   end
19: end

20: function  $\mathbf{u} = \text{SMOOTH}(A, \mathbf{u}, \mathbf{f})$ 
21:   for  $\nu$  steps do
22:      $\mathbf{u} = \mathbf{u} + Q^{-1}(\mathbf{f} - A\mathbf{u})$ 
23:   end
24: end

```

3. Low-rank approximation. In this section we consider a technique designed to reduce computational effort, in terms of both time and memory use, using low-rank methods. We begin with the observation that the solution vector of the Galerkin system (2.7)

$$\mathbf{u} = [u_{11}, u_{21}, \dots, u_{N_x 1}, \dots, u_{1N_\xi}, u_{2N_\xi}, \dots, u_{N_x N_\xi}]^T \in \mathbb{R}^{N_x N_\xi}$$

can be restructured as a matrix

$$(3.1) \quad U = \text{mat}(\mathbf{u}) = \begin{pmatrix} u_{11} & u_{12} & \cdots & u_{1N_\xi} \\ u_{21} & u_{22} & \cdots & u_{2N_\xi} \\ \vdots & \vdots & \ddots & \vdots \\ u_{N_x 1} & u_{N_x 2} & \cdots & u_{N_x N_\xi} \end{pmatrix} \in \mathbb{R}^{N_x \times N_\xi}.$$

Then (2.7) is equivalent to a system in matrix format,

$$(3.2) \quad \mathcal{A}(U) = F,$$

where

$$(3.3) \quad \begin{aligned} \mathcal{A}(U) &= K_0 U G_0^T + \sum_{l=1}^m K_l U G_l^T, \\ F &= \text{mat}(\mathbf{f}) = \text{mat}(g_0 \otimes f_0) = f_0 g_0^T. \end{aligned}$$

It has been shown in [3, 9] that the “matricized” version of the solution U can be well approximated by a low-rank matrix when $N_x N_\xi$ is large. Evidence of this can be seen in Figure 3.1, which shows the singular values of the exact solution U for the benchmark problem discussed in section 4. In particular, the singular values decay exponentially, and low-rank approximate solutions can be obtained by dropping terms from the singular value decomposition corresponding to small singular values.




FIG. 3.1. *Decay of singular values of solution matrix U . Left: exponential covariance, $b = 5$, $h = 2^{-6}$, $m = 8$, $p = 3$. Right: squared exponential covariance, $b = 2$, $h = 2^{-6}$, $m = 3$, $p = 3$. See the benchmark problem in section 4.*

Now we use low-rank approximation in the multigrid solver for (3.2). Let $U^{(i)} = \text{mat}(\mathbf{u}^{(i)})$ be the i th iterate, expressed in matricized format¹, and suppose $U^{(i)}$ is represented as the outer product of two rank- k matrices, i.e., $U^{(i)} \approx V^{(i)}W^{(i)T}$, where $V^{(i)} \in \mathbb{R}^{N_x \times k}$, $W^{(i)} \in \mathbb{R}^{N_\xi \times k}$. This factored form is convenient for implementation and can be readily used in basic matrix operations. For instance, the sum of two matrices gives

$$(3.4) \quad V_1^{(i)}W_1^{(i)T} + V_2^{(i)}W_2^{(i)T} = [V_1^{(i)}, V_2^{(i)}][W_1^{(i)}, W_2^{(i)}]^T.$$

Similarly, $\mathcal{A}(V^{(i)}W^{(i)T})$ can also be written as an outer product of two matrices:

$$(3.5) \quad \begin{aligned} \mathcal{A}(V^{(i)}W^{(i)T}) &= (K_0V^{(i)})(G_0W^{(i)})^T + \sum_{l=1}^m (K_lV^{(i)})(G_lW^{(i)})^T \\ &= [K_0V^{(i)}, K_1V^{(i)}, \dots, K_mV^{(i)}][G_0W^{(i)}, G_1W^{(i)}, \dots, G_mW^{(i)}]^T. \end{aligned}$$

If $V^{(i)}, W^{(i)}$ are used to represent iterates in the multigrid solver and $k \ll \min(N_x, N_\xi)$, then both memory and computational (matrix-vector products) costs can be reduced, from $N_x N_\xi$ to $(N_x + N_\xi)k$. Note, however, that the ranks of the iterates may grow due to matrix additions. For example, in (3.5) the rank may increase from k to $(m+1)k$ in the worst case. A way to prevent this from happening, and also to keep costs low, is to truncate the iterates and force their ranks to remain low.

¹In the sequel, we use $\mathbf{u}^{(i)}$ and $U^{(i)}$ interchangeably to represent the equivalent vectorized or matricized quantities.

3.1. Low-rank truncation. Our truncation strategy is derived using an idea from [9]. Assume $\tilde{X} = \tilde{V}\tilde{W}^T$, $\tilde{V} \in \mathbb{R}^{N_x \times \tilde{k}}$, $\tilde{W} \in \mathbb{R}^{N_\xi \times \tilde{k}}$, and $X = \mathcal{T}(\tilde{X})$ is truncated to rank k with $X = VW^T$, $V \in \mathbb{R}^{N_x \times k}$, $W \in \mathbb{R}^{N_\xi \times k}$ and $k < \tilde{k}$. First, compute the QR factorization for both \tilde{V} and \tilde{W} ,

$$(3.6) \quad \tilde{V} = Q_{\tilde{V}}R_{\tilde{V}}, \quad \tilde{W} = Q_{\tilde{W}}R_{\tilde{W}}, \quad \text{so } \tilde{X} = Q_{\tilde{V}}R_{\tilde{V}}^T Q_{\tilde{W}}^T R_{\tilde{W}}^T.$$

The matrices $R_{\tilde{V}}$ and $R_{\tilde{W}}$ are of size $\tilde{k} \times \tilde{k}$. Next, compute a singular value decomposition (SVD) of the small matrix $R_{\tilde{V}}R_{\tilde{W}}^T$:

$$(3.7) \quad R_{\tilde{V}}R_{\tilde{W}}^T = \hat{V}\text{diag}(\sigma_1, \dots, \sigma_{\tilde{k}})\hat{W}^T$$

where $\sigma_1, \dots, \sigma_{\tilde{k}}$ are the singular values in descending order. We can truncate to a rank- k matrix where k is specified using either a relative criterion for singular values,

$$(3.8) \quad \sqrt{\sigma_{k+1}^2 + \dots + \sigma_{\tilde{k}}^2} \leq \epsilon_{\text{rel}} \sqrt{\sigma_1^2 + \dots + \sigma_{\tilde{k}}^2}$$

or an absolute one,

$$(3.9) \quad k = \max\{k \mid \sigma_k \geq \epsilon_{\text{abs}}\}.$$

Then the truncated matrices can be written in MATLAB notation as

$$(3.10) \quad V = Q_{\tilde{V}}\hat{V}(:, 1:k), \quad W = Q_{\tilde{W}}\hat{W}(:, 1:k)\text{diag}(\sigma_1, \dots, \sigma_k).$$

Note that the low-rank matrices X obtained from (3.8) and (3.9) satisfy

$$(3.11) \quad \|X - \tilde{X}\|_F \leq \epsilon_{\text{rel}}\|\tilde{X}\|_F$$

and

$$(3.12) \quad \|X - \tilde{X}\|_F \leq \epsilon_{\text{abs}}\sqrt{\tilde{k} - k},$$

respectively. The right-hand side of (3.12) is bounded by $\sqrt{N_\xi}\epsilon_{\text{abs}}$ since in general $N_\xi < N_x$. The total cost of this computation is $O((N_x + N_\xi + \tilde{k})\tilde{k}^2)$, which is much cheaper than computing an SVD for X directly.

3.2. Low-rank multigrid. The multigrid solver with low-rank truncation is given in Algorithm 3.1. It uses truncation operators \mathcal{T}_{rel} and \mathcal{T}_{abs} , which are defined using a relative and an absolute criterion, respectively. In each iteration, one multigrid cycle (VCYCLE) is applied to the residual equation. Since the overall magnitudes of the singular values of the correction matrix $C^{(i)}$ decrease as $U^{(i)}$ converges to the exact solution (see Figure 3.2 for example), it is suitable to use a relative truncation tolerance ϵ_{rel} inside the VCYCLE function. In the smoothing function (SMOOTH), the iterate is truncated after each smoothing step using a relative criterion

$$(3.13) \quad \|\mathcal{T}_{\text{rel}_1}(U) - U\|_F \leq \epsilon_{\text{rel}}\|F^h - \mathcal{A}^h(U_0^h)\|_F$$

where A^h , U_0^h , and F^h are arguments of the VCYCLE function, and $F^h - \mathcal{A}^h(U_0^h)$ is the residual at the beginning of each V-cycle. In Line 13, the residual is truncated via a more stringent relative criterion

$$(3.14) \quad \|\mathcal{T}_{\text{rel}_2}(R^h) - R^h\|_F \leq \epsilon_{\text{rel}}h\|F^h - \mathcal{A}^h(U_0^h)\|_F$$

Algorithm 3.1: Multigrid with low-rank truncation

```

1: initialization:  $i = 0$ ,  $R^{(0)} = F$  in low-rank format,  $r_0 = \|F\|_F$ 
2: while  $r > tol * r_0$  &  $i \leq maxit$  do
3:    $C^{(i)} = \text{VCYCLE}(A, 0, R^{(i)})$ 
4:    $\tilde{U}^{(i+1)} = U^{(i)} + C^{(i)}$ ,  $U^{(i+1)} = \mathcal{T}_{\text{abs}}(\tilde{U}^{(i+1)})$ 
5:    $\tilde{R}^{(i+1)} = F - \mathcal{A}(U^{(i+1)})$ ,  $R^{(i+1)} = \mathcal{T}_{\text{abs}}(\tilde{R}^{(i+1)})$ 
6:    $r = \|R^{(i+1)}\|_F$ ,  $i = i + 1$ 
7: end

8: function  $U^h = \text{VCYCLE}(A^h, U_0^h, F^h)$ 
9:   if  $h == h_0$  then
10:    solve  $\mathcal{A}^h(U^h) = F^h$  directly
11:   else
12:     $U^h = \text{SMOOTH}(A^h, U_0^h, F^h)$ 
13:     $\tilde{R}^h = F^h - \mathcal{A}^h(U^h)$ ,  $R^h = \mathcal{T}_{\text{rel}_2}(\tilde{R}^h)$ 
14:     $R^{2h} = \mathcal{R}(R^h)$ 
15:     $C^{2h} = \text{VCYCLE}(A^{2h}, 0, R^{2h})$ 
16:     $U^h = U^h + \mathcal{P}(C^{2h})$ 
17:     $U^h = \text{SMOOTH}(A^h, U^h, F^h)$ 
18:   end
19: end

20: function  $U = \text{SMOOTH}(A, U, F)$ 
21:   for  $\nu$  steps do
22:     $\tilde{U} = U + \mathcal{S}(F - \mathcal{A}(U))$ ,  $U = \mathcal{T}_{\text{rel}_1}(\tilde{U})$ 
23:   end
24: end

```




FIG. 3.2. Singular values of the correction matrix $C^{(i)}$ at multigrid iteration $i = 0, 1, \dots, 5$ (without truncation) for the benchmark problem in subsection 4.1 (with $\sigma = 0.01$, $b = 5$, $h = 2^{-6}$, $m = 8$, $p = 3$).

where h is the mesh size. In the main **while** loop, an absolute truncation criterion (3.9) with tolerance ϵ_{abs} is used and all the singular values of $U^{(i)}$ below ϵ_{abs} are dropped. The algorithm is terminated either when the largest singular value of the residual matrix $R^{(i)}$ is smaller than ϵ_{abs} or when the multigrid solution reaches the specified accuracy.

Note that the post-smoothing is not explicitly required in Algorithms 2.1 and 3.1, and we include it just for sake of completeness. Also, in Algorithm 3.1, if the smoothing operator has the form $\mathcal{S} = S_1 \otimes S_2$, then for any matrix with a low-rank factorization $X = VW^T$, application of the smoothing operator gives

$$(3.15) \quad \mathcal{S}(X) = \mathcal{S}(VW^T) = (S_2V)(S_1W)^T,$$

so that the result is again the outer product of two matrices of the same low rank. The prolongation and restriction operators (2.13) are implemented in a similar manner. Thus, the smoothing and grid-transfer operators do not affect the ranks of matricized quantities in Algorithm 3.1.

3.3. Convergence analysis. In order to show that Algorithm 3.1 is convergent, we need to know how truncation affects the contraction of error. Consider the case of a two-grid algorithm for the linear system $A\mathbf{u} = \mathbf{f}$, where the coarse-grid solve is exact and no post-smoothing is done. Let \bar{A} be the coefficient matrix on the coarse grid, let $\mathbf{e}^{(i)} = \mathbf{u} - \mathbf{u}^{(i)}$ be the error associated with $\mathbf{u}^{(i)}$, and let $\mathbf{r}^{(i)} = \mathbf{f} - A\mathbf{u}^{(i)} = A\mathbf{e}^{(i)}$ be the residual. It is shown in [4] that if no truncation is done, the error after a two-grid cycle becomes

$$(3.16) \quad \mathbf{e}_{\text{notrunc}}^{(i+1)} = (A^{-1} - \mathcal{S}\bar{A}^{-1}\mathcal{R})A(I - Q^{-1}A)^\nu \mathbf{e}^{(i)},$$

and

$$(3.17) \quad \|\mathbf{e}_{\text{notrunc}}^{(i+1)}\|_A \leq C\eta(\nu)\|\mathbf{e}^{(i)}\|_A,$$

where ν is the number of pre-smoothing steps, C is a constant, and $\eta(\nu) \rightarrow 0$ as $\nu \rightarrow \infty$. The proof consists of establishing the smoothing property

$$(3.18) \quad \|A(I - Q^{-1}A)^\nu \mathbf{y}\|_2 \leq \eta(\nu)\|\mathbf{y}\|_A, \quad \forall \mathbf{y} \in \mathbb{R}^{N_x N_\xi},$$

and the approximation property

$$(3.19) \quad \|(A^{-1} - \mathcal{S}\bar{A}^{-1}\mathcal{R})\mathbf{y}\|_A \leq C\|\mathbf{y}\|_2, \quad \forall \mathbf{y} \in \mathbb{R}^{N_x N_\xi},$$

and applying these bounds to (3.16).

Now we derive an error bound for Algorithm 3.1. The result is presented in two steps. First, we consider the VCycle function only; the following lemma shows the effect of the relative truncations defined in (3.13) and (3.14).

LEMMA 3.1. *Let $\mathbf{u}^{(i+1)} = \text{VCycle}(A, \mathbf{u}^{(i)}, \mathbf{f})$ and let $\mathbf{e}^{(i+1)} = \mathbf{u} - \mathbf{u}^{(i+1)}$ be the associated error. Then*

$$(3.20) \quad \|\mathbf{e}^{(i+1)}\|_A \leq C_1(\nu)\|\mathbf{e}^{(i)}\|_A,$$

where, for small enough ϵ_{rel} and large enough ν , $C_1(\nu) < 1$ independent of the mesh size h .

Proof. For $s = 1, \dots, \nu$, let $\tilde{\mathbf{u}}_s^{(i)}$ be the quantity computed after application of the smoothing operator at step s before truncation, and let $\mathbf{u}_s^{(i)}$ be the modification obtained from truncation by $\mathcal{T}_{\text{rel}_1}$ of (3.13). For example,

$$(3.21) \quad \tilde{\mathbf{u}}_1^{(i)} = \mathbf{u}^{(i)} + Q^{-1}(\mathbf{f} - A\mathbf{u}^{(i)}), \quad \mathbf{u}_1^{(i)} = \mathcal{T}_{\text{rel}_1}(\tilde{\mathbf{u}}_1^{(i)}).$$

Denote the associated error as $\mathbf{e}_s^{(i)} = \mathbf{u} - \mathbf{u}_s^{(i)}$. From (3.13), we have

$$(3.22) \quad \mathbf{e}_1^{(i)} = (I - Q^{-1}A)\mathbf{e}^{(i)} + \delta_1^{(i)}, \quad \text{where } \|\delta_1^{(i)}\|_2 \leq \epsilon_{\text{rel}}\|\mathbf{r}^{(i)}\|_2.$$

Similarly, after ν smoothing steps,

$$(3.23) \quad \begin{aligned} \mathbf{e}_\nu^{(i)} &= (I - Q^{-1}A)^\nu \mathbf{e}^{(i)} + \Delta_\nu^{(i)} \\ &= (I - Q^{-1}A)^\nu \mathbf{e}^{(i)} + (I - Q^{-1}A)^{\nu-1} \delta_1^{(i)} + \dots + (I - Q^{-1}A) \delta_{\nu-1}^{(i)} + \delta_\nu^{(i)}, \end{aligned}$$

where

$$(3.24) \quad \|\delta_s^{(i)}\|_2 \leq \epsilon_{\text{rel}}\|\mathbf{r}^{(i)}\|_2, \quad s = 1, \dots, \nu.$$

In Line 13 of Algorithm 3.1, the residual $\tilde{\mathbf{r}}_\nu^{(i)} = A\mathbf{e}_\nu^{(i)}$ is truncated to $\mathbf{r}_\nu^{(i)}$ via (3.14), so that

$$(3.25) \quad \|\mathbf{r}_\nu^{(i)} - \tilde{\mathbf{r}}_\nu^{(i)}\|_2 \leq \epsilon_{\text{rel}}h\|\mathbf{r}^{(i)}\|_2.$$

Let $\tau^{(i)} = \mathbf{r}_\nu^{(i)} - \tilde{\mathbf{r}}_\nu^{(i)}$. Referring to (3.16) and (3.23), we can write the error associated with $\mathbf{u}^{(i+1)}$ as

$$(3.26) \quad \begin{aligned} \mathbf{e}^{(i+1)} &= \mathbf{e}_\nu^{(i)} - \mathcal{P}\bar{A}^{-1}\mathcal{R}\mathbf{r}_\nu^{(i)} \\ &= (I - \mathcal{P}\bar{A}^{-1}\mathcal{R}A)\mathbf{e}_\nu^{(i)} - \mathcal{P}\bar{A}^{-1}\mathcal{R}\tau^{(i)} \\ &= \mathbf{e}_{\text{notrunc}}^{(i+1)} + (A^{-1} - \mathcal{P}\bar{A}^{-1}\mathcal{R})A\Delta_\nu^{(i)} - \mathcal{P}\bar{A}^{-1}\mathcal{R}\tau^{(i)} \\ &= \mathbf{e}_{\text{notrunc}}^{(i+1)} + (A^{-1} - \mathcal{P}\bar{A}^{-1}\mathcal{R})(A\Delta_\nu^{(i)} + \tau^{(i)}) - A^{-1}\tau^{(i)}. \end{aligned}$$

Applying the approximation property (3.19) gives

$$(3.27) \quad \|(A^{-1} - \mathcal{P}\bar{A}^{-1}\mathcal{R})(A\Delta_\nu^{(i)} + \tau^{(i)})\|_A \leq C(\|A\Delta_\nu^{(i)}\|_2 + \|\tau^{(i)}\|_2).$$

Using the fact that for any matrix $B \in \mathbb{R}^{N_x N_\xi \times N_x N_\xi}$,

$$(3.28) \quad \sup_{\mathbf{y} \neq \mathbf{0}} \frac{\|B\mathbf{y}\|_A}{\|\mathbf{y}\|_A} = \sup_{\mathbf{y} \neq \mathbf{0}} \frac{\|A^{1/2}B\mathbf{y}\|_2}{\|A^{1/2}\mathbf{y}\|_2} = \sup_{\mathbf{z} \neq \mathbf{0}} \frac{\|A^{1/2}BA^{-1/2}\mathbf{z}\|_2}{\|\mathbf{z}\|_2} = \|A^{1/2}BA^{-1/2}\|_2,$$

we get

$$(3.29) \quad \begin{aligned} \|A(I - Q^{-1}A)^{\nu-s}\delta_s^{(i)}\|_2 &\leq \|A^{1/2}\|_2 \|(I - Q^{-1}A)^{\nu-s}\delta_s^{(i)}\|_A \\ &\leq \|A^{1/2}\|_2 \|A^{1/2}(I - Q^{-1}A)^{\nu-s}A^{-1/2}\|_2 \|\delta_s^{(i)}\|_A \\ &\leq \rho(I - Q^{-1}A)^{\nu-s} \|A^{1/2}\|_2^2 \|\delta_s^{(i)}\|_2 \end{aligned}$$

where ρ is the spectral radius. We have used the fact that $A^{1/2}(I - Q^{-1}A)^{\nu-s}A^{-1/2}$ is a symmetric matrix (assuming Q is symmetric). Define $d_1(\nu) = (\rho(I - Q^{-1}A)^{\nu-1} + \dots + \rho(I - Q^{-1}A) + 1)\|A^{1/2}\|_2^2$. Then (3.24) and (3.25) imply that

$$(3.30) \quad \begin{aligned} \|A\Delta_\nu^{(i)}\|_2 + \|\tau^{(i)}\|_2 &\leq \epsilon_{\text{rel}}(d_1(\nu) + h)\|\mathbf{r}^{(i)}\|_2 \\ &\leq \epsilon_{\text{rel}}(d_1(\nu) + h)\|A^{1/2}\|_2 \|\mathbf{e}^{(i)}\|_A. \end{aligned}$$

On the other hand,

$$\begin{aligned}
 \|A^{-1}\tau^{(i)}\|_A &= (A^{-1}\tau^{(i)}, \tau^{(i)})^{1/2} \leq \|A^{-1}\|_2^{1/2} \|\tau^{(i)}\|_2 \\
 (3.31) \qquad \qquad \qquad &\leq \epsilon_{\text{rel}} h \|A^{-1}\|_2^{1/2} \|\mathbf{r}^{(i)}\|_2 \\
 &\leq \epsilon_{\text{rel}} h \|A^{-1}\|_2^{1/2} \|A^{1/2}\|_2 \|\mathbf{e}^{(i)}\|_A.
 \end{aligned}$$

Combining (3.17), (3.26), (3.27), (3.30), and (3.31), we conclude that

$$(3.32) \qquad \qquad \qquad \|\mathbf{e}^{(i+1)}\|_A \leq C_1(\nu) \|\mathbf{e}^{(i)}\|_A$$

where

$$(3.33) \qquad C_1(\nu) = C\eta(\nu) + \epsilon_{\text{rel}}(C(d_1(\nu) + h) + h\|A^{-1}\|_2^{1/2})\|A^{1/2}\|_2.$$

Note that $\rho(I - Q^{-1}A) < 1$, $\|A\|_2$ is bounded by a constant, and $\|A^{-1}\|_2$ is of order $O(h^{-2})$ [14]. Thus, for small enough ϵ_{rel} and large enough ν , $C_1(\nu)$ is bounded below 1 independent of h . \square

Next, we adjust this argument by considering the effect of the absolute truncations in the main **while** loop. In Algorithm 3.1, the **VCYCLE** is used for the residual equation, and the updated solution $\tilde{\mathbf{u}}^{(i+1)}$ and residual $\tilde{\mathbf{r}}^{(i+1)}$ are truncated to $\mathbf{u}^{(i+1)}$ and $\mathbf{r}^{(i+1)}$, respectively, using an absolute truncation criterion as in (3.9). Thus, at the i th iteration ($i > 1$), the residual passed to the **VCYCLE** function is in fact a perturbed residual, i.e.,

$$(3.34) \qquad \mathbf{r}^{(i)} = \tilde{\mathbf{r}}^{(i)} + \beta = A\mathbf{e}^{(i)} + \beta, \quad \text{where } \|\beta\|_2 \leq \sqrt{N_\xi} \epsilon_{\text{abs}}.$$

It follows that in the first smoothing step,

$$(3.35) \qquad \tilde{\mathbf{u}}_1^{(i)} = \mathbf{u}^{(i)} + Q^{-1}(\mathbf{f} - A\mathbf{u}^{(i)} + \beta), \quad \mathbf{u}_1^{(i)} = \mathcal{T}_{\text{rel}_1}(\tilde{\mathbf{u}}_1^{(i)}),$$

and this introduces an extra term in $\Delta_\nu^{(i)}$ (see (3.23)),

$$(3.36) \quad \Delta_\nu^{(i)} = (I - Q^{-1}A)^{\nu-1} \delta_1^{(i)} + \dots + (I - Q^{-1}A) \delta_{\nu-1}^{(i)} + \delta_\nu^{(i)} - (I - Q^{-1}A)^{\nu-1} Q^{-1} \beta.$$

As in the derivation of (3.29), we have

$$(3.37) \qquad \|A(I - Q^{-1}A)^{\nu-1} Q^{-1} \beta\|_2 \leq \rho(I - Q^{-1}A)^{\nu-1} \|A^{1/2}\|_2^2 \|Q^{-1}\|_2 \|\beta\|_2.$$

In the case of a damped Jacobi smoother (see (4.4)), $\|Q^{-1}\|_2$ is bounded by a constant. Denote $d_2(\nu) = \rho(I - Q^{-1}A)^{\nu-1} \|A^{1/2}\|_2^2 \|Q^{-1}\|_2$. Also note that $\|\mathbf{r}^{(i)}\|_2 \leq \|A^{1/2}\|_2 \|\mathbf{e}^{(i)}\|_A + \|\beta\|_2$. Then (3.30) and (3.31) are modified to

$$\begin{aligned}
 (3.38) \qquad &\|A\Delta_\nu^{(i)}\|_2 + \|\tau^{(i)}\|_2 \\
 &\leq \epsilon_{\text{rel}}(d_1(\nu) + h) \|\mathbf{r}^{(i)}\|_2 + d_2(\nu) \|\beta\|_2 \\
 &\leq \epsilon_{\text{rel}}(d_1(\nu) + h) \|A^{1/2}\|_2 \|\mathbf{e}^{(i)}\|_A + (d_2(\nu) + \epsilon_{\text{rel}}(d_1(\nu) + h)) \|\beta\|_2,
 \end{aligned}$$

and

$$(3.39) \qquad \|A^{-1}\tau^{(i)}\|_A \leq \epsilon_{\text{rel}} h \|A^{-1}\|_2^{1/2} \|A^{1/2}\|_2 \|\mathbf{e}^{(i)}\|_A + \epsilon_{\text{rel}} h \|A^{-1}\|_2^{1/2} \|\beta\|_2.$$

As we truncate the updated solution $\tilde{\mathbf{u}}^{(i+1)}$, we have

$$(3.40) \qquad \mathbf{u}^{(i+1)} = \tilde{\mathbf{u}}^{(i+1)} + \gamma, \quad \text{where } \|\gamma\|_2 \leq \sqrt{N_\xi} \epsilon_{\text{abs}}.$$

Let

$$(3.41) \quad C_2(\nu) = Cd_2(\nu) + \epsilon_{\text{rel}}(C(d_1(\nu) + h) + h\|A^{-1}\|_2^{1/2}) + \|A^{1/2}\|_2.$$

From (3.38)–(3.41), we conclude with the following theorem:

THEOREM 3.2. *Let $\mathbf{e}^{(i)} = \mathbf{u} - \mathbf{u}^{(i)}$ denote the error at the i th iteration of Algorithm 3.1. Then*

$$(3.42) \quad \|\mathbf{e}^{(i+1)}\|_A \leq C_1(\nu)\|\mathbf{e}^{(i)}\|_A + C_2(\nu)\sqrt{N_\xi}\epsilon_{\text{abs}},$$

where $C_1(\nu) < 1$ for large enough ν and small enough ϵ_{rel} , and $C_2(\nu)$ is bounded by a constant. Also, (3.42) implies that

$$(3.43) \quad \|\mathbf{e}^{(i)}\|_A \leq C_1^i(\nu)\|\mathbf{e}^{(0)}\|_A + \frac{1 - C_1^i(\nu)}{1 - C_1(\nu)}C_2(\nu)\sqrt{N_\xi}\epsilon_{\text{abs}},$$

i.e., the A -norm of the error for the low-rank multigrid solution at the i th iteration is bounded by $C_1^i(\nu)\|\mathbf{e}^{(0)}\|_A + O(\sqrt{N_\xi}\epsilon_{\text{abs}})$. Thus, Algorithm 3.1 converges until the A -norm of the error becomes as small as $O(\sqrt{N_\xi}\epsilon_{\text{abs}})$.

It can be shown that the same result holds if post-smoothing is used. Also, the convergence of full (recursive) multigrid with these truncation operations can be established following an inductive argument analogous to that in the deterministic case (see, e.g., [5, 8]). Besides, in Algorithm 3.1, the truncation on $\tilde{\mathbf{r}}^{(i+1)}$ imposes a stopping criterion, i.e.,

$$(3.44) \quad \begin{aligned} \|\tilde{\mathbf{r}}^{(i+1)}\|_2 &\leq \|\tilde{\mathbf{r}}^{(i+1)} - \mathbf{r}^{(i+1)}\|_2 + \|\mathbf{r}^{(i+1)}\|_2 \\ &\leq \sqrt{N_\xi}\epsilon_{\text{abs}} + \text{tol} * r_0. \end{aligned}$$

In section 4 we will vary the value of ϵ_{abs} and see how the low-rank multigrid solver works compared with Algorithm 2.1 where no truncation is done.

Remark 3.3. If instead a relative truncation is used in the **while** loop so that

$$(3.45) \quad \mathbf{r}^{(i+1)} = \tilde{\mathbf{r}}^{(i+1)} + \beta = A\mathbf{e}^{(i+1)} + \beta, \quad \text{where } \|\beta\|_2 \leq \epsilon_{\text{rel}}\|\tilde{\mathbf{r}}^{(i+1)}\|,$$

then a similar convergence result can be derived, and the algorithm stops when

$$(3.46) \quad \|\tilde{\mathbf{r}}^{(i+1)}\|_2 \leq \frac{\text{tol} * r}{1 - \epsilon_{\text{rel}}}.$$

However, the relative truncation in general results in a larger rank for $\mathbf{r}^{(i)}$, and the improvement in efficiency will be less significant.

4. Numerical experiments. Consider the benchmark problem with a two-dimensional spatial domain $D = (-1, 1)^2$ and constant source term $f = 1$. We look at two different forms for the covariance function $r(x, y)$ of the diffusion coefficient $c(x, \omega)$.

4.1. Exponential covariance. The exponential covariance function takes the form

$$(4.1) \quad r(x, y) = \sigma^2 \exp\left(-\frac{1}{b}\|x - y\|_1\right).$$

This is a convenient choice because there are known analytic solutions for the eigenpair $(\lambda_l, c_l(x))$ [7]. In the KL expansion, take $c_0(x) = 1$ and $\{\xi_l\}_{l=1}^m$ independent and uniformly distributed on $[-1, 1]$:

$$(4.2) \quad c(x, \omega) = c_0(x) + \sqrt{3} \sum_{l=1}^m \sqrt{\lambda_l} c_l(x) \xi_l(\omega).$$

Then $\sqrt{3}\xi_l$ has zero mean and unit variance, and Legendre polynomials are used as basis functions for the stochastic space. The correlation length b affects the decay of $\{\lambda_l\}$ in the KL expansion. The number of random variables m is chosen so that

$$(4.3) \quad \left(\sum_{l=1}^m \lambda_l \right) / \left(\sum_{l=1}^M \lambda_l \right) \geq 95\%.$$

Here M is a large number which we set as 1000.

We now examine the performance of the multigrid solver with low-rank truncation. We employ a damped Jacobi smoother, with

$$(4.4) \quad Q = \frac{1}{\omega} D, \quad D = \text{diag}(A) = I \otimes \text{diag}(K_0),$$

and apply three smoothing steps ($\nu = 3$) in the SMOOTH function. Set the multigrid $\text{tol} = 10^{-6}$. As shown in (3.44), the relative residual $\|F - \mathcal{A}(U^{(i)})\|_F / \|F\|_F$ for the solution $U^{(i)}$ produced in Algorithm 3.1 is related to the value of the truncation tolerance ϵ_{abs} . In all the experiments, we also run the multigrid solver without truncation to reach a relative residual that is closest to what we get from the low-rank multigrid solver. We fix the relative truncation tolerance ϵ_{rel} as 10^{-2} . (The truncation criteria in (3.13) and (3.14) are needed for the analysis. In practice we found the performance with the relative criterion in (3.11) to be essentially the same as the results shown in this section.) The numerical results, i.e., the rank of multigrid solution, the number of iterations, and the elapsed time (in seconds) for solving the Galerkin system, are given in Tables 4.1 to 4.3. In all the tables, the 3rd and 4th columns are the results of low-rank multigrid with different values of truncation tolerance ϵ_{abs} , and for comparison the last two columns show the results for the multigrid solver without truncation. The Galerkin systems are generated from the Incompressible Flow and Iterative Solver Software (IFISS, [16]). All computations are done in MATLAB 9.0.0 (R2016a) on a MacBook with 1.6 GHz Intel Core i5 and 4 GB SDRAM.

Table 4.1 shows the performance of the multigrid solver for various mesh sizes h , or spatial degrees of freedom N_x , with other parameters fixed. The 3rd and 5th columns show that multigrid with low-rank truncation uses less time than the standard multigrid solver. This is especially true when N_x is large: for $h = 2^{-8}$, $N_x = 261121$, low-rank approximation reduces the computing time from 2857s to 347s. The improvement is much more significant (see the 4th and 6th columns) if the problem does not require very high accuracy for the solution. Table 4.2 shows the results for various degrees of freedom N_ξ in the stochastic space. The multigrid solver with absolute truncation tolerance 10^{-6} is more efficient compared with no truncation in all cases. The 4th and 6th columns indicate that the decrease in computing time by low-rank truncation is more obvious with the larger tolerance 10^{-4} .

We have observed that when the standard deviation σ in the covariance function (4.1) is smaller, the singular values of the solution matrix U decay faster (see Figure 3.1), and it is more suitable for low-rank approximation. This is also shown in

TABLE 4.1

Performance of multigrid solver with $\epsilon_{abs} = 10^{-6}$, 10^{-4} , and no truncation for various $N_x = (2/h - 1)^2$. Exponential covariance, $\sigma = 0.01$, $b = 4$, $m = 11$, $p = 3$, $N_\xi = 364$.

		$\epsilon_{abs} = 10^{-6}$	$\epsilon_{abs} = 10^{-4}$	No truncation	
64×64 grid $h = 2^{-5}$ $N_x = 3969$	Rank	51	12		
	Iterations	5	4	5	4
	Elapsed time	7.55	1.65	12.60	10.08
	Rel residual	1.51e-6	6.05e-5	9.97e-7	1.38e-5
128×128 grid $h = 2^{-6}$ $N_x = 16129$	Rank	51	12		
	Iterations	6	4	5	3
	Elapsed time	22.55	5.32	54.59	32.92
	Rel residual	2.45e-6	9.85e-5	1.23e-6	2.20e-4
256×256 grid $h = 2^{-7}$ $N_x = 65025$	Rank	49	13		
	Iterations	5	4	5	3
	Elapsed time	77.72	24.59	311.27	188.70
	Rel residual	4.47e-6	2.07e-4	1.36e-6	2.35e-04
512×512 grid $h = 2^{-8}$ $N_x = 261121$	Rank	39	16		
	Iterations	5	3	4	3
	Elapsed time	347.24	89.52	2857.82	2099.06
	Rel residual	9.93e-6	4.33e-4	1.85e-5	2.43e-4

TABLE 4.2

Performance of multigrid solver with $\epsilon_{abs} = 10^{-6}$, 10^{-4} , and no truncation for various $N_\xi = (m+p)!/(m!p!)$. Exponential covariance, $\sigma = 0.01$, $h = 2^{-6}$, $p = 3$, $N_x = 16129$.

		$\epsilon_{abs} = 10^{-6}$	$\epsilon_{abs} = 10^{-4}$	No truncation	
$b = 5, m = 8$ $N_\xi = 165$	Rank	25	9		
	Iterations	5	4	5	3
	Elapsed time	6.73	1.75	19.33	11.65
	Rel residual	5.06e-6	3.41e-4	1.22e-6	2.20e-4
$b = 4, m = 11$ $N_\xi = 364$	Rank	51	12		
	Iterations	6	4	5	3
	Elapsed time	22.55	5.32	54.59	32.92
	Rel residual	2.45e-6	9.85e-5	1.23e-6	2.20e-4
$b = 3, m = 16$ $N_\xi = 969$	Rank	91	23		
	Iterations	6	5	5	4
	Elapsed time	119.53	18.02	197.82	158.56
	Rel residual	5.71e-7	3.99e-5	1.23e-6	1.63e-5
$b = 2.5, m = 22$ $N_\xi = 2300$	Rank	165	86		
	Iterations	8	5	6	4
	Elapsed time	915.92	187.69	1033.29	682.45
	Rel residual	1.59e-7	8.57e-6	9.29e-8	1.63e-5

the numerical results. In the previous cases, we fixed σ as 0.01. In Table 4.3, the advantage of low-rank multigrid is clearer for a smaller σ , and the solution is well approximated by a matrix of smaller rank. On the other hand, as we increase the value of σ , the rank of the solution matrix and the computing time also increase. This is mainly due to the overhead introduced in the truncation operators when the ranks of the iterates are relatively large. In particular, the QR factorization is the most

expensive part in the low-rank multigrid solver and it takes more than half of the total time used for solving the linear system (Table 4.4). Another observation from the above numerical experiments is that the iteration counts are largely unaffected by truncation. In Algorithm 3.1, similar numbers of iterations are required to reach a comparable accuracy as in the cases with no truncation.

TABLE 4.3

Performance of multigrid solver with $\epsilon_{abs} = 10^{-6}$, 10^{-4} , and no truncation for various σ . Exponential covariance, $b = 4$, $h = 2^{-6}$, $m = 11$, $p = 3$, $N_x = 16129$, $N_\xi = 364$.

		$\epsilon_{abs} = 10^{-6}$	$\epsilon_{abs} = 10^{-4}$	No truncation	
$\sigma = 0.001$	Rank	13	12		
	Iterations	6	4	5	4
	Elapsed time	7.93	3.85	54.43	43.58
	Rel residual	1.09e-6	6.53e-5	1.22e-6	1.63e-5
$\sigma = 0.01$	Rank	51	12		
	Iterations	6	4	5	3
	Elapsed time	22.55	5.32	54.59	32.92
	Rel residual	2.45e-6	9.85e-5	1.23e-6	2.20e-4
$\sigma = 0.1$	Rank	136	54		
	Iterations	6	4	5	3
	Elapsed time	129.82	22.33	55.49	33.62
	Rel residual	3.28e-6	2.47e-4	1.88e-6	2.62e-4
$\sigma = 0.3$	Rank	234	128		
	Iterations	9	7	8	4
	Elapsed time	508.81	130.86	86.77	43.42
	Rel residual	6.03e-6	4.71e-4	2.99e-6	7.76e-4

TABLE 4.4

Time used for QR factorization and SVD in Algorithm 3.1. Exponential covariance, $b = 4$, $h = 2^{-6}$, $m = 11$, $p = 3$, $N_x = 16129$, $N_\xi = 364$, $\epsilon_{abs} = 10^{-6}(10^{-4})$.

σ	0.01	0.1	0.3
QR	61.0%(50.9%)	73.6%(62.8%)	78.4%(72.1%)
SVD	8.4%(9.4%)	5.7%(8.4%)	4.8%(6.1%)

4.2. Squared exponential covariance. In the second example we consider covariance function

$$(4.5) \quad r(x, y) = \sigma^2 \exp \left(-\frac{1}{b^2} \|x - y\|_2^2 \right).$$

The eigenpair $(\lambda_l, c_l(x))$ is computed via a Galerkin approximation of the eigenvalue problem

$$(4.6) \quad \int_D r(x, y) c_l(y) dy = \lambda_l c_l(x).$$

Again, in the KL expansion (4.2), take $c_0(x) = 1$ and $\{\xi_l\}_{l=1}^m$ independent and uniformly distributed on $[-1, 1]$. The eigenvalues of the squared exponential covariance (4.5) decay much faster than those of (4.1), and thus fewer terms are required to satisfy (4.3). For instance, for $b = 2$, $m = 3$ will suffice. Table 4.5 shows the performance

of multigrid with low-rank truncation for various spatial degrees of freedom N_x . In this case, we are able to work with finer meshes since the value of N_ξ is smaller. In all experiments the low-rank multigrid solver uses less time compared with no truncation.

TABLE 4.5
Performance of multigrid solver with $\epsilon_{abs} = 10^{-6}$, 10^{-4} , and no truncation for various $N_x = (2/h - 1)^2$. Squared exponential covariance, $\sigma = 0.01$, $b = 2$, $m = 3$, $p = 3$, $N_\xi = 20$.

		$\epsilon_{abs} = 10^{-6}$	$\epsilon_{abs} = 10^{-4}$	No truncation	
128 \times 128 grid $h = 2^{-6}$ $N_x = 16129$	Rank	9	4		
	Iterations	5	3	4	3
	Elapsed time	0.84	0.34	1.08	0.82
	Rel residual	1.20e-5	9.15e-4	1.63e-5	2.20e-4
256 \times 256 grid $h = 2^{-7}$ $N_x = 65025$	Rank	8	4		
	Iterations	4	3	4	3
	Elapsed time	2.69	1.41	4.58	3.46
	Rel residual	3.99e-5	9.09e-4	1.78e-05	2.35e-4
512 \times 512 grid $h = 2^{-8}$ $N_x = 261121$	Rank	8	2		
	Iterations	4	2	4	2
	Elapsed time	10.68	2.28	18.93	9.61
	Rel residual	6.41e-5	6.91e-3	1.85e-5	3.29e-3
1024 \times 1024 grid $h = 2^{-9}$ $N_x = 1045629$	Rank	8	2		
	Iterations	4	2	4	2
	Elapsed time	54.28	10.82	115.75	63.54
	Rel residual	6.41e-5	6.93e-3	1.90e-5	3.32e-3

5. Conclusions. In this work we focused on the multigrid solver, one of the most efficient iterative solvers, for the stochastic steady-state diffusion problem. We discussed how to combine the idea of low-rank approximation with multigrid to reduce computational costs. We proved the convergence of the low-rank multigrid method with an analytic error bound. It was shown in numerical experiments that the low-rank truncation is useful in decreasing the computing time when the variance of the random coefficient is relatively small. The proposed algorithm also exhibited great advantage for problems with large number of spatial degrees of freedom.

REFERENCES

- [1] I. BABUŠKA, R. TEMPONE, AND G. E. ZOURARIS, *Galerkin finite element approximations of stochastic elliptic partial differential equations*, SIAM Journal on Numerical Analysis, 42 (2004), pp. 800–825.
- [2] J. BALLANI AND L. GRASEDYCK, *A projection method to solve linear systems in tensor format*, Numerical Linear Algebra with Applications, 20 (2013), pp. 27–43.
- [3] P. BENNER, A. ONWUNTA, AND M. STOLL, *Low-rank solution of unsteady diffusion equations with stochastic coefficients*, SIAM/ASA Journal on Uncertainty Quantification, 3 (2015), pp. 622–649.
- [4] H. ELMAN AND D. FURNIVAL, *Solving the stochastic steady-state diffusion problem using multigrid*, IMA Journal of Numerical Analysis, 27 (2007), pp. 675–688.
- [5] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*, Oxford University Press, UK, second ed., 2014.
- [6] O. G. ERNST AND E. ULLMANN, *Stochastic Galerkin matrices*, SIAM Journal on Matrix Analysis and Applications, 31 (2010), pp. 1848–1872.

- [7] R. G. GHANEM AND P. D. SPANOS, *Stochastic Finite Elements: A Spectral Approach*, Dover Publications, New York, 2003.
- [8] W. HACKBUSCH, *Multi-Grid Methods and Applications*, Springer-Verlag Berlin Heidelberg, 1985.
- [9] D. KRESSNER AND C. TOBLER, *Low-rank tensor Krylov subspace methods for parametrized linear systems*, SIAM Journal of Matrix Analysis and Applications, 32 (2011), pp. 1288–1316.
- [10] K. LEE AND H. C. ELMAN, *A preconditioned low-rank projection method with a rank-reduction scheme for stochastic partial differential equations*, May 2016, arXiv:1605.05297 [math.NA].
- [11] M. LOÈVE, *Probability Theory*, Van Nostrand, New York, 1960.
- [12] G. J. LORD, C. E. POWELL, AND T. SHARDLOW, *An Introduction to Computational Stochastic PDEs*, Cambridge University Press, 2014.
- [13] H. G. MATTHIES AND E. ZANDER, *Solving stochastic systems with low-rank tensor compression*, Linear Algebra and its Applications, 436 (2012), pp. 3819–3838.
- [14] C. E. POWELL AND H. C. ELMAN, *Block-diagonal preconditioning for spectral stochastic finite-element systems*, IMA Journal of Numerical Analysis, 29 (2009), pp. 350–375.
- [15] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, SIAM, second ed., 2003.
- [16] D. SILVESTER, H. ELMAN, AND A. RAMAGE, *Incompressible Flow and Iterative Solver Software (IFISS)*, Aug. 2015, <http://www.manchester.ac.uk/ifiss> (accessed 2016/05/27). Version 3.4.
- [17] D. XIU AND G. E. KARNIADAKIS, *Modeling uncertainty in flow simulations via generalized polynomial chaos*, Journal of Computational Physics, 187 (2003), pp. 137–167.